




# The Linux cluster of the H1 experiment at DESY

Ralf Gerhards  
Bogdan Lobodzinski  
H1 Collaboration  
DESY

Large-Scale Cluster Computing Workshop  
Fermilab, May 22 – 25, 2001

Ralf Gerhards, DESY FH1

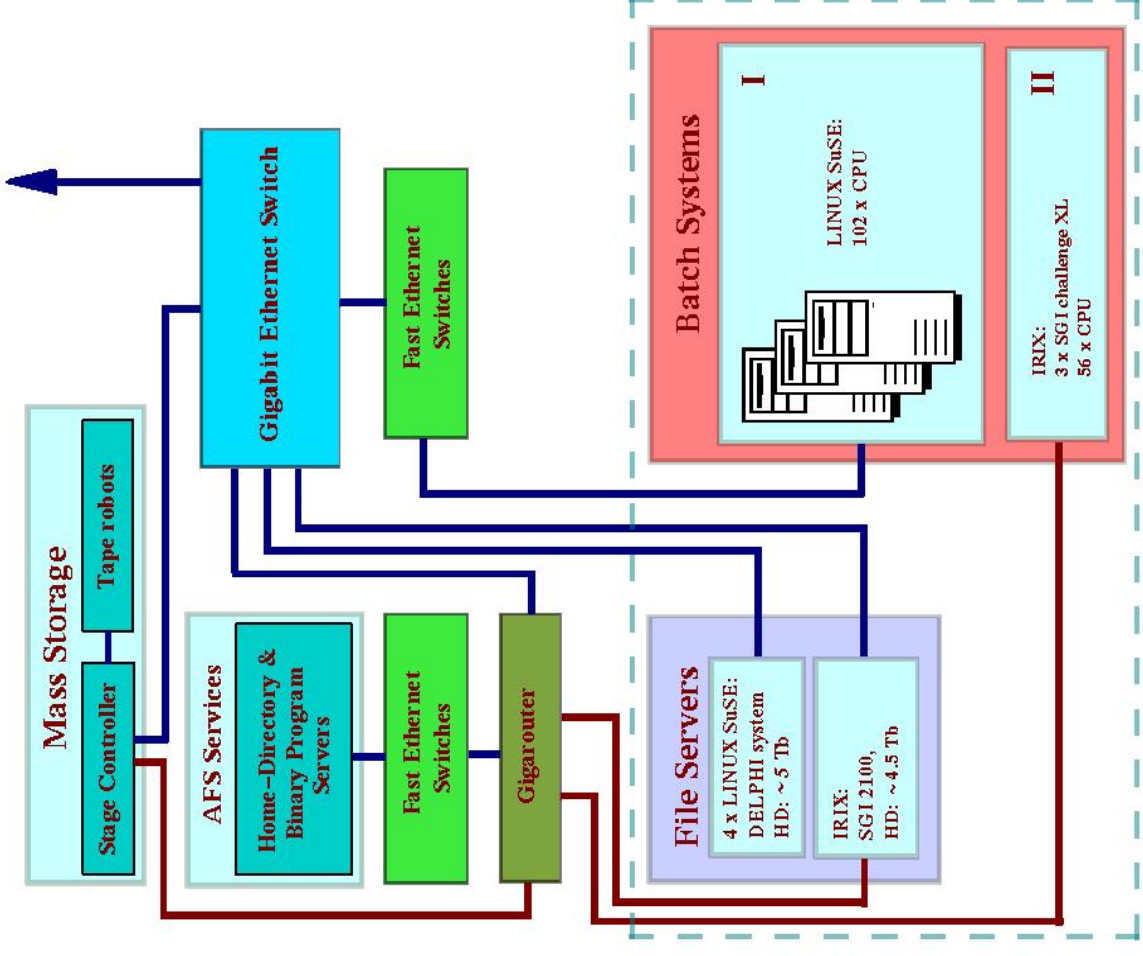
LCCWS, Fermilab, May 22, 2001



# Introduction

- multi-purpose computing environment for
  - Higher Level Trigger (L45)
  - Reprocessing
  - Monte Carlo Production
  - Batch Analysis
- assignment of nodes to these tasks very dynamic
- similar farms for ZEUS, HERMES, HERA-B, ..

# Setup




# Hardware Components

- **~50 processing nodes**
  - dual processor (PII/PIII)
  - some local disk (SCSI/IDE)
  - fastEthernet
- **disk servers**
  - SGI Origin 2001 (~5 TB, mostly fibre channel)
  - 4 DELFI systems (~5 TB, IDE)
  - access via RFIO
- **tape servers**
  - 2 Powderhorn silos
  - 1 Grau silo
  - access via OSM
  - access via intermediate disk cache in progress
- **networking**
  - fastEthernet between processing nodes
  - 10–20 nodes form a subfarm
  - uplink via GbitEthernet
  - part of (SGI) servers via Hippi

# Linux File Servers

- 2 PIII 1 GHZ
- 3 IDE controller
- 22 75GB IDE disks
- GBitEthernet



Ralf Gerhards, DESY FH1



LCCWS, Fermilab, May 22, 2001

# Software Components

- **Batch system**
  - PBS
  - open source
  - integrated with AFS
- **Data access**
  - AFS
  - RFIO
  - Disk Cache
  - transparent access to all data on tape
  - integrated with I/O system (FPAACK)
- **H1 framework (for event distribution)**
  - based on Corba
  - initially for high level trigger, reprocessing
  - eventually for batch analyses as well

# Installation

- (automatic) common installation by DESY IT staff
- using SuSE Linux and AFS
- common for farms and desktops
- central services, based on Solaris servers
- **production system**
- DESY Linux 3
  - SuSE 6.3
  - glibc 2.1.2
  - 2.2.x kernel (x = 18, 19)
- farms at DESY
  - H1
    - ca. 50 nodes
  - ZEUS
    - ca. 60 nodes
  - Hermes
    - ca. 20 nodes
  - Hera-B
    - ca. 50 nodes
    - plus trigger farms
- **12 Linux file servers**

# Conclusions

- fits very well with Linux based desktop environment
- very well received by users (old SGI batch servers hardly used anymore)
- investigate new data storage models
  - move processes to data
  - use Linux file servers
- profit from common installation and monitoring service by DESY/IT